arXiv:1603.00663v2 [cs.RO] 3 Sep 2016

Unsupervised Watertight Mesh Generation for Physics Simulation Applications Using Growing Neural Gas on Noisy Free-Form Object Models

Tobias Fromm, Christian A. Mueller, and Andreas Birk *

September 6, 2016

Abstract

We present a framework to generate watertight mesh representations in an unsupervised manner from noisy point clouds of complex, heterogeneous objects with free-form surfaces. The resulting meshes are ready to use in applications like kinematics and dynamics simulation where watertightness and fast processing are the main quality criteria. This works with no necessity of user interaction, mainly by utilizing a modified Growing Neural Gas technique for surface reconstruction combined with several post-processing steps.

In contrast to existing methods, the proposed framework is able to cope with input point clouds generated by consumer-grade RGBD sensors and works even if the input data features large holes, e.g. a missing bottom which was not covered by the sensor. Additionally, we explain a method to unsupervisedly optimize the parameters of our framework in order to improve generalization quality and, at the same time, keep the resulting meshes as coherent as possible to the original object regarding visual and geometric properties.

1 Introduction

Surface reconstruction from point clouds has often been addressed in computer graphics, robotics, medical computing and other fields. However, depending on sampling device (e.g. consumer-grade sensors like the Kinect) and application, input points may be too sparse and noisy for conventional methods [1] [2] [3] [4] to deliver appropriately polygonized surfaces. Additionally, for many applications like physics simulations, watertight meshes (featuring no holes) and even volume meshes (in addition to the surface, featuring structures also inside the mesh) are mandatory.

We present a framework which, in contrast to existing methods, is able to deliver a watertight volume mesh for any free-form [5] object, if necessary, on the expense of a roughly approximated surface. This works not only for geometric primitives, but also for more complex, heterogeneous objects with free-form surfaces. The proposed method bases on our previous work which deals with surface reconstruction using *Growing Neural Gas (GNG)* [6] and the registration of noisy 2.5D point clouds from different viewpoints to create a 3D object point cloud [7]. Although this point cloud generally misses the bottom of the objects due to self-occlusion, we employ several established methods to still create meshes usable in physics simulations.

The generated volume mesh representations can be used, for instance, for simulation purposes within a robotic manipulation scenario where they are shown in a virtual scene and a virtual robot can interact with them. Our main application lies in the EU project *RobLog* with the objective to automate logistic processes, such as unloading containers [8]. The used hardware setup and the

*Tobias Fromm and Christian A. Mueller contributed equally to this work. All authors are with the Robotics Group, Computer Science & Electrical Engineering, Jacobs University Bremen, Germany; *[t.fromm,chr.mueller]@jacobs-university.de*. This research has received funding from the European Union's Seventh Framework Programme (EU FP7 ICT-2) within the project "Cognitive Robot for the Automation of Logistic Processes (RobLog)".

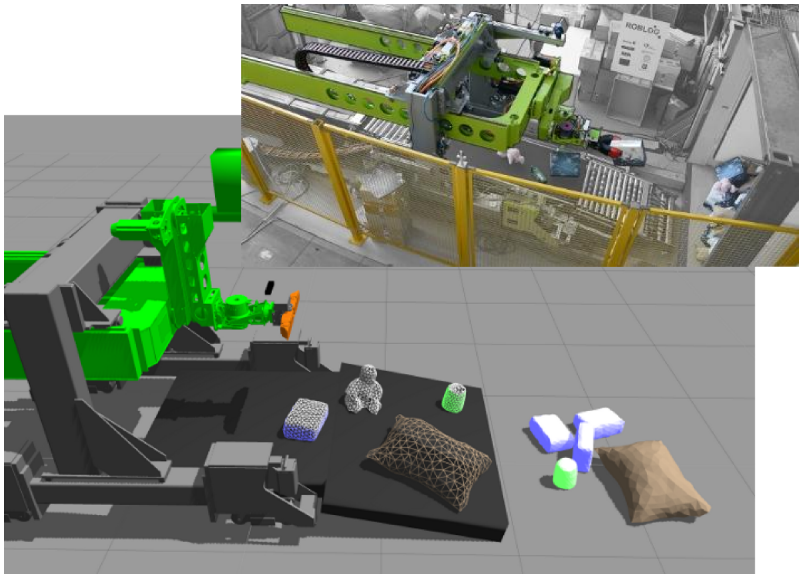

Figure 1: Application scenario: simulation of real-world object handling

corresponding *Gazebo* [9] physics simulation scene are shown in Fig. 1. Due to the nature of physics engines, watertightness is an essential requirement which prevents models from interpenetrating.

Additionally, in order to improve the generalization performance as well as the accuracy of our approach, we present a method to optimize its parameters in an unsupervised manner using *Particle Swarm Optimization (PSO)* [10].

The dataset we use in this paper is publicly available[1] and provides raw sensor data, registered point clouds and the generated meshes along with Gazebo models.

The contribution of our work is:

- a framework which requires no user-provided parameters and no user interaction that
- combines several established methods to generate watertight object meshes ready to use in physics simulations
- from noisy RGBD point clouds of complex objects
- including an unsupervised technique to optimize essential parameters.

To the authors' best knowledge, our framework is the first one of its kind regarding the entirety of the above criteria.

2 Related Work

As for surface reconstruction, several methods exist which have been used frequently on different kinds of reconstruction problems. We use them as a baseline for our method, as described in Section 5.2.

One class of these methods bases on determining a function of the surface given by the input point set and subsequently creating polygons from the surface function using Marching Cubes [11]. The first one of these is Hoppe et al.'s method [2] which estimates a signed distance function and determines its zero set. Secondly, Carr et al.'s approach [3] uses a *Radial Basis Function (RBF)* to describe the input data and again polygonizes the result via Marching Cubes.

Equally well-known is the method of Kazhdan et al. [1] which is based on optimizing Poisson equations for approximating a surface. This works similarly to using RBFs, but adds hierarchies in the equations which allows for locally fine-grained reconstruction where necessary. On the other

[1] http://robotics.jacobs-university.de/datasets/2016-watertight-meshes-v01

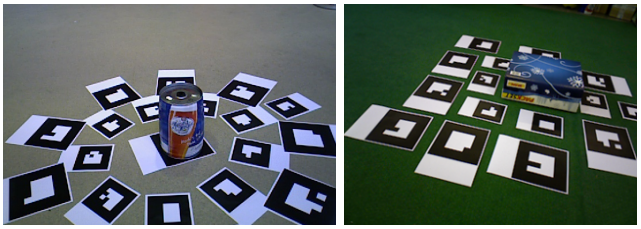

Figure 2: Object modeling setup used in our previous work [7]

hand, RBFs, working globally on the input data, may prove impractical on input data like ours with varying density.

Another method we took as a baseline was Kolluri et al.'s *Eigencrust* [4] for which its authors emphasize its particular noise-robustness. It uses a Delaunay tetrahedralization of the input points of which each tetrahedron is labelled as *inside* or *outside* the object. Subsequently, the triangles which intersect between either set are used as the requested surface. Their main contribution in comparison to previous approaches is the usage of a spectral partitioning technique to divide the point sets which they claim to be more effective in distinguishing the object's edges and the environment.

Additionally, we already extensively evaluated *KinectFusion* [12] in our previous work [7]. Hence, it will not be taken into account again in this paper since it has already been discussed in detail that, due to some deficiencies, *KinectFusion* do not provide models of the needed quality.

In the context of our application, we found the mentioned methods to deliver insufficient results, as we will explain in detail in the following sections. For this reason, the work described in this paper uses *Growing Neural Gas* which interprets surface reconstruction as a learning problem. Growing Neural Gas which was proposed by Fritzke [13] is an unsupervised learning technique that can be applied to learn the topology of a given distribution. The learned topology is reflected by a graphical representation consisting of vertices (so-called neurons) and edges. We exploit GNG learning to learn the topology of a given point cloud which can be seen as a three-dimensional distribution. The resulting graph can be interpreted as a surface mesh and used as such for further processing steps. In our previous work [6] we adapted the original GNG approach [13] to surface reconstruction-specific properties. Using GNG for surface reconstruction shows several advantages as described in Section 3.2, amongst others, *online learning*, which is the ability to stop the reconstruction process any time while always receiving a triangulated surface.

Online learning has, for instance, been used in recent work of Vierjahn and Hinrichs [14] which introduces a reconstruction method comparable to our approach, but with no guarantee for watertightness which is a crucial requirement in our desired application scenario.

Growing Neural Gas relies on some parameters which change its behavior in several ways. Leaving these parameters on default values as given by Fritzke [13] and our previous work [6] yields non-optimal results because GNG is not able to adapt its topology to the desired granularity due to different requirements defined in their work. In contrast to other approaches which deal with surface reconstruction and use a fixed, heuristically-determined parameter set [14] [15] [16], we provide a methodology to determine the quality of the generated GNG representation and to optimize its parameters in an unsupervised manner. In previous work about Growing Neural Gas, this problem has not been tackled yet either [17] or motivated in a very specific environment where it is hard to predict its general behavior [18].

Since the GNG-generated meshes from our method still feature some deficiencies, like missing triangulation if holes in the input data exceed a certain size, we propose several post-processing steps, among others, hole filling. Thorough investigation has been performed on this topic already by many authors [19] [20] [21], so we decided to apply a shown-to-work method, namely *MeshFix* [22], instead of displacing our focus in this direction.

Our framework, similar to all other surface reconstruction methods mentioned above, relies on multi-view 3D data. Since we want to create watertight models which cannot be penetrated by other objects, a full 3D model is required. 2.5D point clouds, as obtained from a 3D sensor from one single viewpoint, provide an incomplete model which misses (self-)occluded parts compared to

the original object's volume. We therefore made use of our previously developed fully-automatic method [7] which can create 3D models without sophisticated equipment, solely using a Kinect-like RGBD sensor and printed Augmented Reality (AR) markers like in Fig. 2. After capturing a number of single-view shots while manually moving the sensor around the object, registration and pose-refinement are performed.

One important feature of our previous work [7] is its insensitivity to variance in environment conditions, e.g. regarding the lighting. Because of its ease of use even for untrained persons, this method fits the environment of the approach presented in this paper where unsupervised processing with no human intervention is desired.

3 Watertight Volume Mesh Generation

In Alg. 1 we provide an overview of the major processing steps of our method after which we will explain every step in detail. In each line in Alg. 1 a reference is given where to find further information beyond of what will be described in the following paragraphs.

Algorithm 1 Mesh Generation

1: **input**: noisy 3D object point cloud, e.g., from [7]
2: reconstruct surface using GNG (→ Alg. 2)
3: remove close-by edges (→ Alg. 3)
4: fill holes using *Meshfix* [22]
5: remove duplicate vertices/faces using *VCG Library* [23]
6: *optional*: simplify mesh using *QSlim* [24]
7: *optional*: tetrahedralize mesh using *TetGen* [25]
8: restore color information (→ Alg. 4)
9: **output**: object volume mesh

3.1 Input Data Quality Requirements

In the presented method, we use object models generated with a Kinect-like low-cost RGBD sensor and AR markers from our previous work [7]. These point cloud models consist of a magnitude of 500k-1.5M points (750k in the example in Fig. 3) including the noise stemming from the consumer-grade sensor and its movement while recording the data for object modeling.

3.2 Surface reconstruction with Growing Neural Gas

Growing Neural Gas is an unsupervised learning technique which aims to learn the topology of a given input distribution. In an iterative manner using competitive Hebbian Learning [13], the topology of the distribution is reflected in a graphical model. Basically, in each iteration a random sample point is selected from the distribution and fired into the graphical model. GNG aims to adapt the graphical model by modifying the number of vertices, edges, and poses of vertices such that the input distribution is optimally approximated. Note that a vertex in the model does not represent a sample point from the input distribution, but a prominent location in the distribution. Since the model is extended during each iteration, this represents an online learning algorithm which can be stopped anytime while always generating a reasonable output distribution.

In case of a three-dimensional input distribution, which in our case is the input point cloud, the finally generated graphical model can be interpreted as a surface mesh. Several beneficial properties can be observed: Due to the iterative manner of adaptation to a three-dimensional input distribution the modified GNG procedure *reorganizes* vertices and edges such that prominent geometric properties are reflected in the resulting mesh. As a consequence, the GNG results in a *denoised* surface mesh – outliers in the input distribution will have a minor or even negligible effect on the surface mesh. Moreover it can be shown that if a sufficient number of iterations is applied, the resulting surface mesh will converge to a *Delaunay-triangulated* mesh [13].

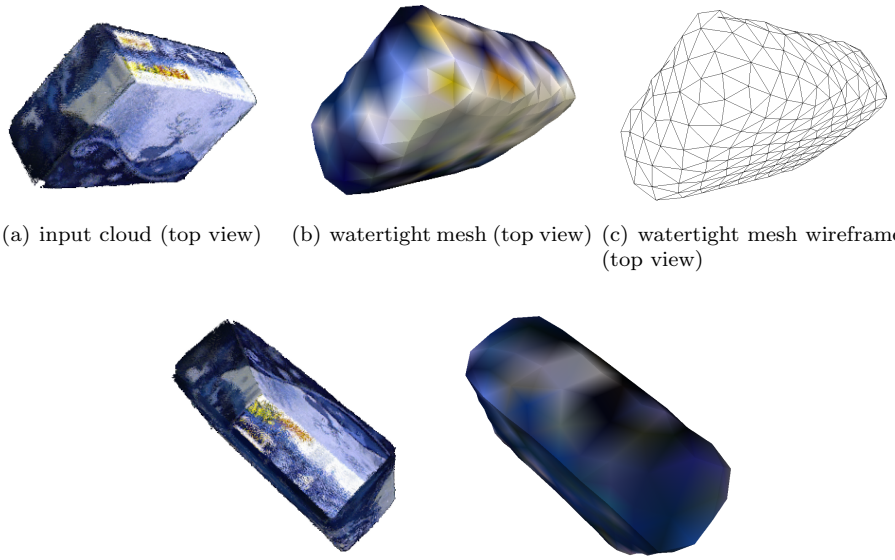

(a) input cloud (top view) (b) watertight mesh (top view) (c) watertight mesh wireframe (top view)

(d) input cloud (bottom view, note the open bottom) (e) watertight mesh (bottom view, hole has been closed)

Figure 3: Watertight mesh generated via the proposed framework with no user interaction

In our previous work [6] we adapted the original GNG approach [13] to the specifics of the surface reconstruction problem. A set of modifications was introduced: (a) instead of adding new vertices in a fixed interval during the learning iterations, we add a new vertex to the model $\mathcal{G}$ only if the accumulated error (Δerror$(\cdot)$ – cf. [13]) of a vertex in $\mathcal{G}$ exceeds a threshold t_γ; (b) we remove diverted vertices which are not coherent with the input distribution; (c) we add mild noise to the input distribution to enhance the triangulation process during training; (d) we perform multiple epochs of retraining on the model to gain a more coherent model relative to the input distribution.

For the last modification an additional convergence criterion *is_converged*$(\cdot)$ (Eq. 1) was introduced to analyze the consistency between the learned GNG model $\mathcal{G}_t$ for training epoch t and the input distribution (point cloud) $\mathcal{P}$,

$$\textit{is_converged}(\mathcal{G}_t) = \begin{cases} \textit{true} \text{ if } \varepsilon(\mathcal{G}_t) > \varepsilon(\mathcal{G}_{t-1}) \\ \textit{false} \text{ otherwise} \end{cases} \tag{1}$$

$$\varepsilon(\mathcal{G}_t) = \frac{\sum_{v \in \mathcal{G}} \arg\min_{p \in \mathcal{P}} (\|v - p\|_2)}{|\mathcal{G}_t|} \tag{2}$$

where $\varepsilon(\mathcal{G}_t)$ is a score which, based on the L^2-*norm*, computes the consistency between $\mathcal{G}_t$ and $\mathcal{P}$, $|\mathcal{G}_t|$ represents the number of vertices in $\mathcal{G}_t$, v denotes a vertex from $\mathcal{G}_t$ and p a point in $\mathcal{P}$. After each retraining epoch we compute *is_converged*$(\mathcal{G}_t)$ for the current state of the retrained GNG model. If *is_converged*$(\mathcal{G}_t)$ returns *true* the training will be stopped and $\mathcal{G}_t$ will represent the final surface mesh for the given point cloud $\mathcal{P}$, otherwise a new training epoch is applied on $\mathcal{G}_t$.

Alg. 2 provides a brief overview of the epochal GNG learning. For further details about the Growing Neural Gas learning procedure we refer to [13] and our previous initial work [6] on which we build here.

3.3 Post-Processing

After surface reconstruction, several post-processing steps are performed to make the mesh watertight, see Alg. 1.

Algorithm 2 GNG-based Surface Reconstruction

1: **input**: point cloud $\mathcal{P}$
2: create an empty GNG Model $\mathcal{G}_t$, $t = 0$
3: **while** $\neg$*is_converged*$(\mathcal{G}_t)$ **do**
4: train GNG on $\mathcal{G}_t$ using $\mathcal{P}$
5: $t = t + 1$
6: **end while**
7: **output**: generated surface mesh $\mathcal{G} = \mathcal{G}_t$

Close-By Edge Removal Edges that are very close to each other, i.e. in the range of the point cloud's resolution, are removed in this step in order to not maintain overlapping triangles. This step, developed by the authors, is described in Alg. 3 and was integrated because the GNG might create polygons of order > 3 with connected vertices. These edges are then very close to each other without necessarily intersecting because the vertices of the created polygons usually do not lie on a flat surface. However, the following post-processing steps cannot deal with such contingencies in a way that a smooth surface is created regardless of the presence of overlapping and intersecting triangles, thus this step is crucial.

Algorithm 3 Close-By Edge Removal

for all vertices $v_i \in \mathcal{G}$ **do**
  for all neighboring vertices v_j of v_i **do**
    for all neighb. vertices v_k of v_j where $v_k \neq v_i$ **do**
      if $d(e_{ij}, e_{jk}) < t_p$ **then**
        remove e_{jk}
      end if
    end for
  end for
end for
where e_{ij} is the edge between vertices v_i and v_j, $d(e_{ij}, e_{jk})$ is the Euclidean distance between the closest points of e_{ij} and e_{jk}, t_p is a threshold for the maximum allowed edge proximity

A suitable value for the threshold t_p in Alg. 3 has to be determined experimentally, depending on the resolution of the input point clouds. The value should be as low as possible in order to not erroneously remove edges of different parts of the model, but still to remove all intersecting or close-by edges. For our input dataset, any value of t_p in the range $[0.005, 0.01]$ m works sufficiently and removes all close-by artifact edges, but leaves the remaining edges intact.

Hole Filling Next, we use *MeshFix* [22] to fill the remaining holes in the mesh. Despite the GNG managing to reconstruct the surface close to the input point cloud, holes may remain due to sparse input data. The reasons for this sparsity may be reflections, illuminance issues, etc. which usually appear in 3D data obtained from low-cost sensors. Additionally, the registered point cloud does not cover the bottom of the object. This leads to a hole which has to be filled. As demonstrated in Fig. 7, this post-processing step enables our framework to deal with big extents of self-occlusion and otherwise missing input data.

Duplicate Face Removal During the hole filling process, duplicate faces might appear. Thus, the next step is to remove them utilizing the respective function of the *VCG Library* [23] because otherwise this may cause the tetrahedralization routine (see below) to be unable to generate tetrahedra between two flatly aligned, congruent faces.

Simplification As an optional step, in case the user desires an even more sparse mesh resolution for faster processing within their application, experiments have shown that *QSlim* [24] provides effective mesh simplification. As shown in Fig. 7, the generated meshes though contain triangles

in a magnitude of several hundreds only and will be coarse-grain enough for sufficient processing speed, so generally this step will not be necessary for most applications.

Tetrahedralization Depending on the application, a tetrahedralization of the resulting surface meshes may be necessary to obtain volume meshes. This is crucial, for instance, for the fidelity of soft-body simulations where an object is not only defined by its boundaries, but also its intrinsic structure. In a dynamics simulation, surface-only meshes will collapse once the physics engine is stepped. Hence, we propose the use of *TetGen* [25] which we successfully use to tetrahedralize the generated surface meshes. For pure rigid-body simulation, this step can be omitted.

Color Information Restauration Eventually, the color information needs to be restored since it was removed during the surface generation and refinement process. Not for all applications it is necessary to have colored meshes, but this is a process which can efficiently be applied. This step makes use of a k-d tree structure and is described in detail in Alg. 4.

Algorithm 4 Color Information Restauration

1: **input**: reconstructed mesh $\mathcal{G}$, colored point cloud $\mathcal{P}$
2: create k-d tree from $\mathcal{P}$
3: **for** all $p_i \in \mathcal{G}$ **do**
4: search tree for (spatial) nearest neighbor of p_i
5: set color of p_i on color of nearest neighbor in $\mathcal{P}$
6: **end for**

4 Mesh quality optimization

Growing Neural Gas requires a set of parameters to be adapted to the quality of the input data in order to create a mesh with properties considered as optimal. On the one hand, high generalization skills during reconstruction are required for noise removal and fast processing, but the generated meshes should not be too coarse to lose details on complex objects. On the other hand, low generalization tends to overfit the input data and does not close holes efficiently.

For these reasons, we present an approach to optimize the GNG parameters as used in Alg. 1 in an unsupervised manner, using several optimization criteria. Both the GNG parameters and mesh optimization criteria will be described in detail in the following.

4.1 Parameters to optimize

In order to achieve well-reconstructed surfaces from noisy point clouds, a set of parameters is required to be appropriately set up. The selected GNG parameters to optimize are shown in Table 1 where $\mathcal{P}$ denotes a three-dimensional distribution (which represents the input point cloud) and $\mathcal{G}$ a trained GNG model on $\mathcal{P}$. These parameters depend on the dimensions of the input point cloud as well as the inherent and encompassing noise and thus need to be optimized.

The procedure to optimize this 6-dimensional parameter set will be described in the following; first, we need to define the criteria with respect to which to optimize.

4.2 Mesh optimization criteria

In the following, we present two criteria we found to be relevant to have an effect on the mesh quality.

Consistency Error $\varepsilon(\mathcal{G})$ defined as the consistency between $\mathcal{G}_t$ and $\mathcal{P}$, see Eq. 2

Table 1: Growing Neural Gas Parameters

parameter	description
ϵ_b	move nearest vertex towards the signal by this fraction
ϵ_n	move all adjacent vertices of the nearest vertex to signal towards the signal by this fraction
$a_{\max}$	remove edges in model $\mathcal{G}$ which are older than $a_{\max}$ within a GNG training epoch
t_γ	a new vertex is added only if the $\Delta\text{error}(\cdot)$ of a vertex $v \in \mathcal{G}$ exceeds t_γ
α	if a new vertex is added decrease the error variable of the nearest and second nearest vertex by this factor
d	factor to decrease the error variable of each vertex in model $\mathcal{G}$

Essential parameters applied in each GNG training epoch (cf. [6] and [13])

Mean Edge Length Ratio $\eta(\mathcal{G})$

$$\eta(\mathcal{G}) = 1 - \frac{\sum_1^n \frac{l_{\min}}{l_{\max}}}{n} \tag{3}$$

where n is the number of triangles in the mesh, $l_{\min}$ is the length of the shortest, $l_{\max}$ the length of the longest edge of the respective triangle.

From these criteria, $\varepsilon(\mathcal{G})$ describes the consistency between the generated mesh and the input data where $\eta(\mathcal{G})$ correlates with the mesh surface appearance.

Both criteria are combined into an *evaluation function* $\theta(\mathcal{G})$ for each time step t in the following way:

$$\theta(\mathcal{G}) = \begin{cases} \textit{true} & \text{if } (\varepsilon_{t+1}(\mathcal{G}) < \varepsilon_t(\mathcal{G}) \wedge \eta_{t+1}(\mathcal{G}) \leq \eta_t(\mathcal{G})) \\ & \quad \vee (\varepsilon_{t+1}(\mathcal{G}) \leq \varepsilon_t(\mathcal{G}) \wedge \eta_{t+1}(\mathcal{G}) < \eta_t(\mathcal{G})) \\ \textit{false} & \text{otherwise} \end{cases} \tag{4}$$

Instead of the condition in Eq. 4, pure *and*-relations, *or*-relations or a product of $\varepsilon(\mathcal{G})$ and $\eta(\mathcal{G})$ would have been possible to use. In our experiments, we disregarded these combinations due to different numeric ranges. Especially since $\varepsilon(\mathcal{G})$ measures a distance, it has no upper boundary; typical values of $\varepsilon(\mathcal{G})$ are in a 10^{-3} magnitude whereas $\eta(\mathcal{G})$ is a ratio of two equivalent values and thus normalized on $[0, 1[$. As for *and* and *or*-relations, one criterion is favored on the expense of the other in case of divergence, thus this will lead to a perturbed optimization.

Another possible mesh optimization criterion is a Hausdorff-like distance between input point clouds and meshes which was proposed to measure surface distances in several works [26] [27]. However, the processing intensity of this measure (an all-to-all comparison of input points and mesh vertices) makes it impractical to use in our case because it increases our framework's runtimes about 10-fold.

4.3 Particle Swarm Optimization

Many existing approaches use hand-crafted values for their parameters to achieve satisfactory results. Moreover, parameters can form strong relationships between each other.

Since the parameter optimization problem presented here is discontinuous, all optimization methods relying on a continuous function describing the problem will fail. As for pattern search-based methods, the brute-force Grid Search is one of the most trivial, but least efficient ones. Instead, we utilize *Particle Swarm Optimization (PSO)* [10] in a parallelized implementation of [28].

PSO can be used without assumptions about the data to optimize. As an input, it is able to deal with a number of numerical parameters in arbitrary ranges (for our case, see Table 1) and needs an evaluation function towards which to optimize ($\theta(\mathcal{G})$). It will first create a random configuration within the 6-dimensional space formed by the GNG parameters and randomly distribute a swarm of agents around it.

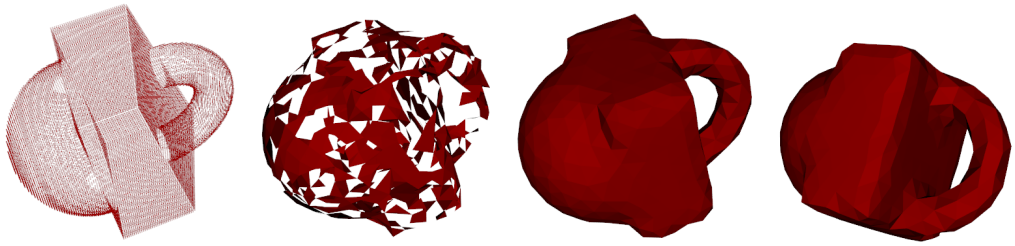

(a) synthetic point cloud used for optimization (b) mesh for initial non-optimized parameters (c) mesh for optimized parameters (view 1) (d) mesh for optimized parameters (view 2)

Figure 4: Particle Swarm Optimization results (visual)

As shown in Alg. 5, with every step, the GNG is trained and the evaluation function $\theta(\mathcal{G})$ is resolved. Afterwards, the center of the agent distribution shifts in the direction of the current global optimum. After a specified maximum number of iterations, the position of the globally optimal agent determines the optimized set of GNG parameters to be used from then on.

Algorithm 5 Particle Swarm Optimization

input: dimensions of parameter space $\mathcal{A}$
initialize all agents $a_i \in \mathcal{A}$ randomly
initialize globally optimal agent $a_{\text{opt}} = a_0$
while less than maximum number of iterations **do**
  for all $a_i \in \mathcal{A}$ **do**
    train GNG using a_i
    if $\theta(\mathcal{G})$ **then**
      update optimal agent $a_{\text{opt}} = a_i$
    end if
  end for
  update all agents' parameter sets to move the swarm towards the current optimum
end while
output: globally optimal agent a_{opt}

We used the default settings for PSO as they yielded satisfying results in our experiments (see Section 5.1). Meta-optimizing the Particle Swarm Optimizer has been subject to research (e.g. [29]), but since the default parameters given by Kennedy and Eberhart [10] provide sufficient optimization convergence for our use case, we regard this topic as out-of-scope for the work proposed in this paper.

5 Results

5.1 Growing Neural Gas parameter optimization

For optimizing the GNG parameters, we introduce a synthetic object model shown in Fig. 4. The rationale behind this is to avoid model-specific bias from our object models. It combines surface features of prototype shapes like sphere, cuboid and torus, containing edges, flat and curved surfaces. This model was intentionally created noise-free because the GNG deals with noise occuring in our real scenario.

PSO was run for 100 iterations using the evaluation function $\theta(\mathcal{G})$, which combines $\eta(\mathcal{G})$ and $\varepsilon(\mathcal{G})$, with the results in Fig. 5. It can be observed that the $\eta(\mathcal{G})$ and $\varepsilon(\mathcal{G})$ scores both improve over time and converge after about 35 iterations, hence this leads to the optimized GNG parameter set given in Table 2.

The respective ranges were pre-set as an initial estimate, yielding faster optimization and reducing the likelihood to fall into local minima which may occur for PSO, especially for discontinous

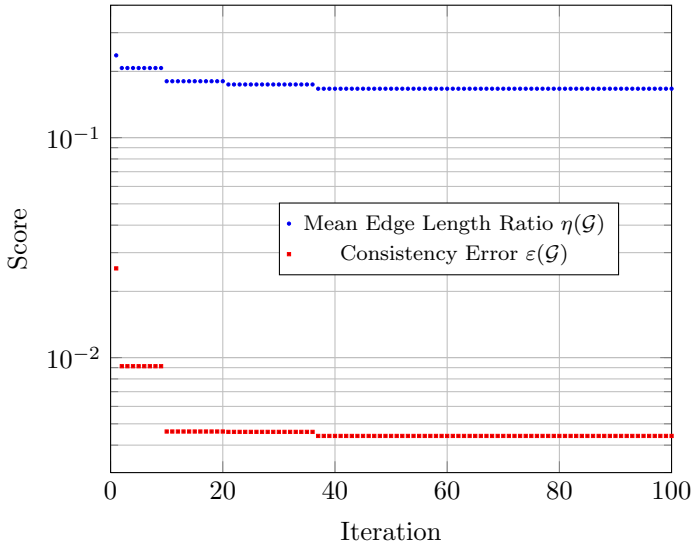

Figure 5: PSO results using $\theta(\mathcal{G})$ as an evaluation function

Table 2: Particle Swarm Optimization results

parameter	**range**	**seed**	**optimized value**
ϵ_b	$[0, 1]$	0.2	0.0739138
ϵ_n	$[0, 0.2]$	0.006	0.00870156
t_γ	$[0, 5]$	3.0	2.72645
$a_{\max}$	$[50, 250]$	60	133
α	$[0, 1]$	0.5	0.521687
d	$[0, 1]$	0.995	0.999321

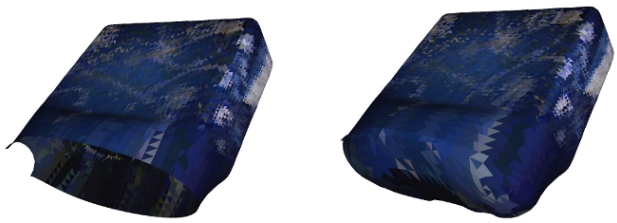

(a) Poisson reconstruction (b) after post-processing

Figure 6: Poisson surface reconstruction results (cf. our results, Fig. 3)

evaluation functions like $\theta(\mathcal{G})$. In order to accelerate the optimization process we initialize the Particle Swarm with seed parameters (cf. Table 2) which are based on the original parameters evaluated by Fritzke [13]. These parameters provide a reasonable guess so that the GNG algorithm is capable to evolve an initial set of vertices.

5.2 Comparison to similar approaches

For qualitatively evaluating our framework, we used *Poisson* [1], *Eigencrust* [4], as well as Hoppe et al.'s [2] and Carr et al.'s [3] methods for baselines which replaced Line 2 in Alg. 1 while not changing the overall process. All of the mentioned methods were evaluated using different objects from our testing set, see the examples in Fig. 7.

As for the reconstruction proximity to the input data, Poisson surface reconstruction delivers visually satisfactory results like in Fig. 6(a) when used with hand-tuned parameters (*maximum tree depth: 6, minimum number of samples per node: 20*). Unfortunately, the method is unable to deal with sharp object edges where it creates bulges instead of preserving the structure of the input data. The proposed post-processing does not help since it fills holes only after the original object structure has been distorted (see Fig. 6(b)). Since our input models miss at least the bottom face, this deficiency invalidates the use of Poisson surface reconstruction for our purpose. Additionally, Poisson reconstruction sometimes creates multiple unconnected meshes for input models containing holes where GNG, by definition, generates exactly one mesh. Hence, we consider Poisson-generated meshes as visually appealing for most objects, but unsatisfactory in the context of realistic simulation.

Equally unsatisfactory is the deployment of any of the other three methods (Hoppe et al.'s and Carr et al.'s, which both make use of Marching Cubes [11], and Eigencrust) on our input data. Carr et al. use dense models with several 100k points which are practically noise-free, Hoppe et al.'s example point sets are not very dense (several 1000), but noise-free by definition since they were sampled from CAD models. Kolluri et al.'s point clouds are noise-free as well from the low-cost sensor point of view, only bearing artifical outliers which are mostly placed far away from the object.

Concretely, Hoppe et al.'s method and *Eigencrust* failed on our input data with zero generated surface vertices regardless of their parameter settings where Carr et al.'s method runs for more than one hour of CPU time without any result.

Summarized, we were unable to create a viable result with any of these methods, probably due to the fact that our input point clouds were created with a low-cost RGBD camera and hence incorporate noise these methods cannot deal with.

5.3 Results for objects of different complexity

The proposed method can not only be applied to objects which may also be modeled by geometric primitives, like the box in Fig. 3, but also for more complex items like barrels, balls, tires, sacks and dolls. Fig. 7 shows objects of different complexity modeled with our approach alongside with the values of their respective optimization criteria, each one featuring incomplete model information in the form of self-occlusion or missing faces. Nevertheless, our framework is capable of resolving these model deficiencies.

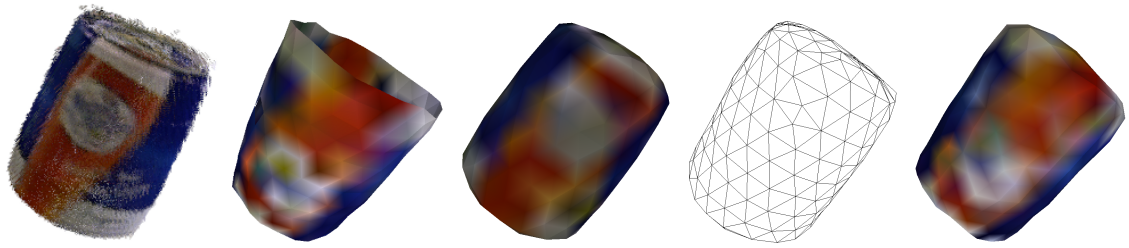

(a) Barrel, $\eta(\mathcal{G}) = 0.209718$, $\varepsilon(\mathcal{G}) = 7.20993 \cdot 10^{-4}$

(b) Coffee sack, $\eta(\mathcal{G}) = 0.204374$, $\varepsilon(\mathcal{G}) = 7.46082 \cdot 10^{-4}$

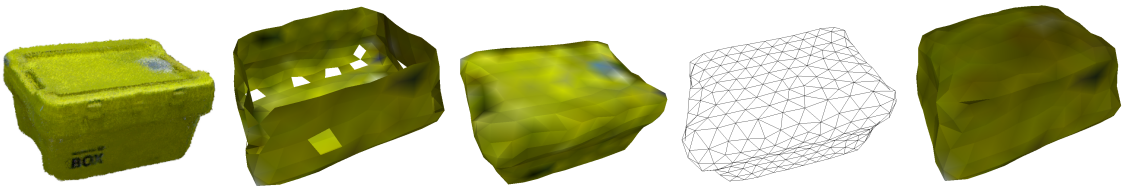

(c) Post box, $\eta(\mathcal{G}) = 0.203831$, $\varepsilon(\mathcal{G}) = 9.24772 \cdot 10^{-4}$

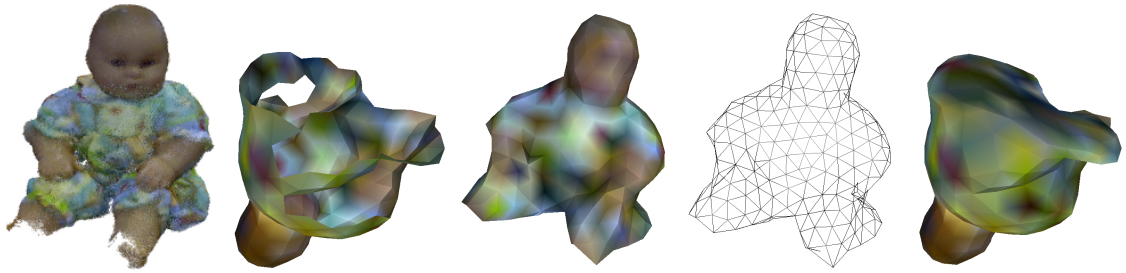

(d) Baby doll, $\eta(\mathcal{G}) = 0.216012$, $\varepsilon(\mathcal{G}) = 8.89707 \cdot 10^{-4}$

Figure 7: Collection of object model point clouds from our previous work [7] and the respective meshes generated with the method proposed herein. Particle Swarm-optimized parameters were used, see Section 5.1.
From left to right: 1. original point cloud, 2. GNG mesh (bottom view), 3. post-processed mesh (top view), 4. post-processed wireframe mesh (top view), 5. post-processed mesh (bottom view).

5.4 Known limitations

In the following, we will identify several cases in which our method may fail. First, since any holes in the meshes will be fixed by a planar series of triangles, no additional information will be generated which is not already present in the input point cloud. In case of the baby doll (Fig. 7(d)–7(d)), not only a planar part of the bottom parallel to ground is missing, but also quite some proportion of its shape due to self-occlusion. Consequently, the reconstructed surface also misses these parts and features a planar surface instead.

The second case where our approach may give non-optimal results when run in an unsupervised way is on fine-detailed objects. An example for this is, again, the baby doll which loses details on its limbs through the generalization performed by the Growing Neural Gas. This is a known limitation of our quality metrics which are unable to determine whether a detail of the model is actually missing or there is noise present. As generally we prefer the automatic generation of watertight models ready to use in simulation, this limitation can be accepted for our application.

6 Conclusion

In this paper, a framework was presented which is capable of automatically generating watertight mesh representations for the use in physics simulations which, in contrary to existing approaches, works in an unsupervised way and can deal with noisy point clouds obtained from a low-cost sensor.

We have shown that our approach works on different kinds of complex objects with free-form surfaces without user interaction, even if not all sides of the object were captured by the sensor. The major parameters of our method have been optimized in an unsupervised manner using a synthetic protoype object in order to generate visually appealing meshes which are close to the input in terms of their geometry. The resulting watertight meshes may boost fidelity in robotic simulation applications.

References

[1] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson Surface Reconstruction," in *Eurographics Symposium on Geometry Processing*, 2006.

[2] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface Reconstruction from Unorganized Points," in *SIGGRAPH*, 1992.

[3] J. Carr, R. Beatson, J. Cherrie, T. Mitchell, W. R. Fright, B. McCallum, and T. Evans, "Reconstruction and representation of 3D objects with radial basis functions," in *SIGGRAPH*, 2001.

[4] R. Kolluri, J. Shewchuk, and J. O'Brien, "Spectral Surface Reconstruction from Noisy Point Clouds," in *SIGGRAPH Symposium on Geometry Processing*, 2004.

[5] R. Campbell and P. Flynn, "A Survey Of Free-Form Object Representation and Recognition Techniques," *Computer Vision and Image Understanding*, vol. 81, no. 2, pp. 166–210, 2001.

[6] C. A. Mueller, N. Hochgeschwender, and P. G. Ploeger, "Surface Reconstruction with Growing Neural Gas," in *IROS 2012 Workshop on Active Semantic Perception*, 2012.

[7] R.-G. Mihalyi, K. Pathak, N. Vaskevicius, T. Fromm, and A. Birk, "Robust 3D Object Modeling with a Low-Cost RGBD-Sensor and AR-Markers for Applications with Untrained End-Users," *Robotics and Autonomous Systems*, vol. 66, pp. 1–17, 2015.

[8] T. Stoyanov, N. Vaskevicius, C. A. Mueller, T. Fromm, R. Krug, V. Tincani, R. Mojtahedzadeh, S. Kunaschk, R. Mortensen Ernits, D. Ricao Canelhas, M. Bonilla, S. Schwertfeger, M. Bonini, H. Halfar, K. Pathak, M. Rohde, G. Fantoni, A. Bicchi, A. Birk, A. Lilienthal, and W. Echelmeyer, "No More Heavy Lifting: Robotic Solutions to the Container Unloading Problem," *Robotics and Automation Magazine*, 2016, to appear.

[9] N. Koenig and A. Howard, "Design and use paradigms for Gazebo, an open-source multi-robot simulator," in *International Conference on Intelligent Robots and Systems*, 2004.

[10] J. Kennedy and R. Eberhart, "Particle Swarm Optimization," in *International Conference on Neural Networks*, 1995.

[11] W. Lorensen and H. Cline, "Marching Cubes: a high resolution 3D surface construction algorithm," *Computer Graphics*, vol. 21, no. 4, pp. 163–169, 1987.

[12] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "KinectFusion: Real-time dense surface mapping and tracking," in *International Symposium on Mixed and Augmented Reality*, 2011.

[13] B. Fritzke, "A Growing Neural Gas Network Learns Topologies," in *Advances in Neural Information Processing Systems*, 1995, pp. 625–632.

[14] T. Vierjahn and K. Hinrichs, "Surface-reconstructing growing neural gas: A method for online construction of textured triangle meshes," *Computers and Graphics*, vol. 51, 2015.

[15] D. Lu, H. Zhao, M. Jiang, S. Zhou, and T. Zhou, "A Surface Reconstruction Method for Highly Noisy Point Clouds," in *Variational, Geometric, and Level Set Methods in Computer Vision*, 2005.

[16] Z. Marton, R. Rusu, and M. Beetz, "On Fast Surface Reconstruction Methods for Large and Noisy Point Clouds," in *International Conference on Robotics and Automation*, 2009.

[17] Y. Holdstein and A. Fischer, "Three-dimensional surface reconstruction using meshing growing neural gas (MGNG)," *The Visual Computer*, vol. 24, no. 4, pp. 295–302, 2008.

[18] J. Garcia-Rodriguez, A. Angelopoulou, J. M. Garcia-Chamizo, A. Psarrou, S. O. Escolano, and V. M. Gimenez, "Autonomous growing neural gas for applications with time constraint: Optimal parameter estimation," *Neural Networks*, vol. 32, pp. 196–208, 2012.

[19] P. Borodin, M. Novotni, and R. Klein, "Progressive Gap Closing for Mesh Repairing," *Science*, 2002.

[20] P. Liepa, "Filling Holes in Meshes," in *Eurographics Symposium on Geometry Processing*, 2003.

[21] T. Ju, "Robust Repair of Polygonal Models," in *SIGGRAPH*, 2004.

[22] M. Attene, "A lightweight approach to repairing digitized polygon meshes," *The Visual Computer*, vol. 26, no. 11, pp. 1393–1406, 2010.

[23] P. Cignoni and F. Ganovelli, "VCG Library," http://vcg.isti.cnr.it/~cignoni/newvcglib/html/.

[24] M. Garland and P. S. Heckbert, "Surface Simplification Using Quadric Error Metrics," in *SIGGRAPH*, 1997.

[25] H. Si and K. Gaertner, "Meshing Piecewise Linear Complexes by Constrained Delaunay Tetrahedralizations," in *International Meshing Roundtable*, 2005.

[26] P. Cignoni, C. Rocchini, and R. Scopigno, "Metro: Measuring Error on Simplified Surfaces," *Computer Graphics Forum*, vol. 17, no. 2, 1998.

[27] N. Aspert, D. Santa-Cruz, and T. Ebrahimi, "MESH: Measuring errors between surfaces using the Hausdorff distance," in *IEEE International Conference on Multimedia and Expo*, 2002.

[28] J. Kennedy, R. Eberhart, and Y. Shi, *Swarm Intelligence*. Morgan Kaufmann Publishers, 2001.

[29] A. Carlisle and G. Dozier, "An Off-The-Shelf PSO," in *Workshop on Particle Swarm Optimization*, 2001.